On Coreferring Text-extracted Event Descriptions with the aid of Ontological Reasoning

Technical Report

Stefano Borgo[1], Loris Bozzato[2], Alessio Palmero Aprosio[2], Marco Rospocher[2], and Luciano Serafini[2]

[1] Laboratory for Applied Ontology, ISTC CNR, Via alla Cascata 56 C, 38123 Trento, Italy
[2] Fondazione Bruno Kessler, Via Sommarive 18, 38123 Trento, Italy

stefano.borgo@cnr.it, {bozzato,aprosio,rospocher,serafini}@fbk.eu

arXiv:1612.00227v1 [cs.AI] 1 Dec 2016

Abstract. Systems for automatic extraction of semantic information about events from large textual resources are now available: these tools are capable to generate RDF datasets about text extracted events and this knowledge can be used to reason over the recognized events.
On the other hand, text based tasks for event recognition, as for example event coreference (i.e. recognizing whether two textual descriptions refer to the same event), do not take into account ontological information of the extracted events in their process.
In this paper, we propose a method to derive event coreference on text extracted event data using semantic based rule reasoning.
We demonstrate our method considering a limited (yet representative) set of event types: we introduce a formal analysis on their ontological properties and, on the base of this, we define a set of coreference criteria. We then implement these criteria as RDF-based reasoning rules to be applied on text extracted event data. We evaluate the effectiveness of our approach over a standard coreference benchmark dataset.

Keywords: Event coreference; Semantic information extraction; Event formalization

1 Introduction

Objects and events are basic ontological categories for knowledge and, thus, for text understanding. They have quite different properties, starting from the way they relate to time and space: objects are primarily related to space while events to time. Furthermore, in reading a book, the book (an object) is wholly present at any time during the reading (an event), while some temporal parts of the reading are not. These examples also show the role of the participation relation between objects and events: an object exists in time by participating in some event (the reader and the book participate in the reading event), and an event is spatially located where its participants are [17].

While research on knowledge extraction from text has traditionally concentrated on objects and their properties. Systems for automatic extraction of semantic information about events are more recent and have largely ignored the events' ontological structure. These tools (e.g., NewsReader [19], Pikes [8]) are clearly important and can be further exploited for other tasks. By generating RDF datasets focused on (text extracted) events, they make possible to use RDF knowledge to reason over the detected events, so that one can now automatically infer new facts. This opportunity allows, for example, to identify conflicting information within and across event knowledge, as well as to discover implicit relationships between events.

As said, text based tasks for event recognition, as for example *event coreference* (i.e. recognizing whether two textual descriptions refer to the same event), do not typically take into account ontological information of the extracted events in their process. In this paper, we aim to show that an ontological analysis of event types [18] contributes positively in this endeavor. To achieve this, we propose a method to derive event coreference on text extracted event data using rule reasoning on the base of a well-founded ontological description of events and their relations. Our proposed method goes as follows:

– We first provide a formal analysis on events and event types by isolating relevant ontological properties.
– On the base of this, we consider a set of criteria to discover (probable) coreferring events.
– We then implement these criteria as RDF-based reasoning rules to be applied on text extracted event data.
– Finally, we evaluate the effectiveness of our approach over a standard coreference benchmark dataset.

In particular, we will exemplify and evaluate our approach on a known benchmark for the evaluation of event coreference, the *EventCorefBank (ECB)*[3] [3] corpus. In presenting each step of our approach, we describe how it has been applied to the events and event types in the corpus and, finally, we show an experimental evaluation of our implementation over the dataset.

Our goal is to demonstrate that, by taking into account the ontological properties of events and their formalization as reasoning rules, it is possible to provide enough flexibility to manage typically large, noisy and incorrect knowledge about text extracted events. We insist that the ontology-based rules provide an additional means for knowledge extraction and should not be considered an alternative to other types of rules.

The main contributions of this paper can be summarized as follows:

– We first propose (in Section 2) a formal analysis of the ontological properties of events and we define our event model.
– In Section 3 we consider the event extraction task and we define a set of tasks (in Section 4) that can be solved via reasoning on the formal representation in our model.
– Then, in Section 5 we consider the particular task of event coreference and we provide its formal characterization.
– Finally, in Section 6 we describe an implementation and evaluation of the approach over the ECB corpus.

2 Events and Event descriptions

We start by stating a formal ontological definition of event, adapted from the approach followed in DOLCE [4]:

> An *Event* is a happening, that is, an entity directly related to time and (indirectly) related to space via its participants [4].

We slightly simplify this view by rephrasing the notion as follows:

> An *Event* identifies a spatio-temporal region, some relevant entities in it, and what happens to these entities.

As usual, we divide the happenings in types depending on the changes that occur to the participants during the time spanned by the event. Formally, we characterize events via four features that we take as ontologically necessary for an event to occur:

1. *Event type* (i.e. the type of happening)
2. *Participant(s)* with their thematic role in the event (i.e. their relevance in the event)
3. *Time of happening*
4. *Global happening location*

For time and location features we might agree on a specified granularity. Note that an event must have at least one participant. We will typically distinguish *agents* (persons, organisms etc.) Vs. *non-agents* (objects and materials) as participants. However, the terms 'active', 'passive' and 'tool' participant are borrowed from the linguistic terminology. For specific types of events, one might also define the "main" participants, i.e. those that characterize and identify an event, and the "accessory" participants.

We use the term 'role' in the ontological sense: the function played by the participant in the event (e.g. a person having the role of public officer in an arrest).

Moreover, we concentrate in particular on *facts*, i.e., events that actually happened (and not, for example, hypothetical happenings or descriptions of events as reported in speeches).

[3] `http://adi.bejan.ro/data/ECB1.0.tar.gz`

2.1 Formal analysis of event types

On the base of this, we can refine the four features by providing a characterization of event types by the following points:

Ontological classification: classification of the event type with respect to DOLCE definition of events (i.e. *accomplishments* or *achievements*) [4].
Active participants: participants seen as active in the event.
Passive participants: participants seen as passive in the event.
Tool participants: participants seen as instruments in the event.
Subevents and sovraevents: types of events that (possibly or necessarily) include or are included in the considered event.
Status before, during and after: conditions on the participants that are verified before, during and after the event.
Relationship among events: (possible or necessary) relations with other events (typically causality).
Symmetrical events: (possible or necessary) events that take place symmetrically with the considered event.
Incompatible events/states: (possible or necessary) events that cannot take place simultaneously with the considered event.
Number of participants: conditions on the number of participants.
Spatial region: conditions on the participants co-location in the considered event.
Temporal region: conditions on the possible duration of the considered event.
Repeatability: conditions on the possibility to repeat an event (w.r.t. the considered participants).

2.2 Event types in ECB

As a first step in using the ECB corpus in our experiments, we applied our analysis on the ECB event types. The *EventCorefBank (ECB)* corpus [3] is a known testset for event coreference. It consists of 482 news texts divided in 43 different topics. A total of 1744 event mentions are annotated, corresponding to 339 distint events. We chose to work on this corpus because it offers a small but relevant set of event types.

We chose 6 of the more representative event (facts) types in the collection (namely: *Arresting, Killing, Dying, Charging, Shooting, Attacking*) and we applied to them our formal analysis. We note that these types of event are closely connected and can give rise to interesting coreference examples. The complete results of the analysis on such event types are provided in Appendix A. For example, in Table 1 we present the analysis of the *Killing* event type. We note that this initial analysis in our approach requires the most effort to be completed[4]: on the other hand, once this effort is completed, the information can be used to define sets of reasoning rules for different tasks and independent from the size and quality of the processed event data. Note that, for some of the event types, their linguistic interpretation (viz. their definition provided in FrameNet) covers multiple types in the ontological classification (e.g., *Arresting* defines both the instantaneous act of seizing someone as well as the whole process of searching, seizing and taking the person into custody). In the current experiment we chose only one of the interpretations (e.g. we consider *Arresting* as an achievement) in order to define an univocal set of coreference rules: however, in a more complete characterization, one may define different analysis for both interpretations and infer from the text context which one is intended.

3 Event extraction from text

Information about events can be gathered from many resources. One of the most commonly used resource for event description is text. News article describes events, hypothetical events, contains event prediction, etc.

[4] We estimated that an expert needs around 20 min. for the analysis of an event type.

Table 1. Formal analysis of event type *Killing*

Ontological classification	Accomplishment
Active participant	Physical object (including agents)
Passive participant	Living entity
Tool participant	Physical object (including chemical and biological entities)
Subevents	Hitting (possible); start of the killing, dying (necessary)
Sovraevents	Murdering, colliding, starving, being infected (all possible)
Status: before	Passive participant is alive
Status: during	Passive participant is alive
Status: after	Passive participant is dead
Relationship among events	Killing causes: damaging
Symmetrical events	Dying for causes external to the passive participant
Incompatible events/state	–
Number of participants	Killing can have one or more active participants; one or more passive participants; one or more tool participants.
Spatial region	Killing happens in the location where the passive participants are located
Temporal region	Interval
Repeatability	There cannot be two killing events with the same passive participant; there can be more than one killing event with the active participant; there can be more than one killing event in the same spatio-temporal region.

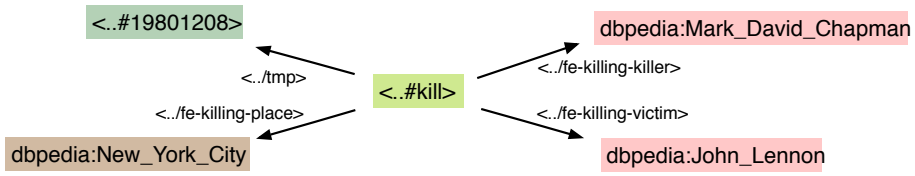

Fig. 1. Graphical representation of the output produced by Pikes on the example sentence.

The growing maturity level of NLP techniques makes now possible to automatic process large amount of textual resources and extract information about the events mentioned in the text. The output however cannot be considered as certain information, as errors might occurs in the extraction process, or wrong information could be present in the source text.

Typically, various types of information about events can be automatically extracted from a text corpus. We review them looking at the output produced by one of the state-of-the-art tools for knowledge extraction, Pikes [8].

Pikes is a Knowledge Extraction framework adopting a 2-phase approach. First—phase 1: *linguistic feature extraction*—an RDF graph of mentions is built by distilling the output of several state-of-the-art NLP tools, including Stanford CoreNLP[5] (tokenization, POS-tagging, lemmatization, NERC, TERN, parsing and coreference resolution), UKB[6] (WSD), DBpedia Spotlight[7] (EL), Mate-tools[8] and Semafor[9] (SRL). Then—phase 2: *knowledge distillation*—the mention graph is processed to distill the knowledge graph using SPARQL-like mapping rules, which are evaluated using RDFpro[10] [7], an RDF manipulation tool used also for RDFS reasoning.

The RDF knowledge graph returned by Pikes contains various typologies of content, that we illustrate using as example the output produced when parsing the sentence "*On December 8, 1980, in NYC, Mark Chapman killed John Lennon.*" (see Figure 1 for a graphical rendering of the output produced or directly access the online demo on Pikes web-site[11]):

[5] `http://nlp.stanford.edu/software/corenlp.shtml`
[6] `http://ixa2.si.ehu.es/ukb/`
[7] `http://spotlight.dbpedia.org/`
[8] `http://code.google.com/p/mate-tools/`
[9] `http://www.cs.cmu.edu/~ark/SEMAFOR/`
[10] `http://rdfpro.fbk.eu/`
[11] `http://pikes.fbk.eu`

– event instances (e.g., "kill"), typed with respect to known ontological (e.g., FrameBase[12], SUMO[13]) and linguistic resources (e.g., FrameNet[14], VerbNet[15]);
– event participants (e.g., "dbpedia:John_Lennon"), i.e., individuals possibly linked to typed instances in linked data repositories (e.g., DBpedia[16]);
– role of participants in the event (e.g., "killing-victim"), according to role catalogs taken from available ontological (e.g., FrameBase) and linguistic resources (e.g., FrameNet, VerbNet);
– temporal information (e.g., "19801208"), grounding the event in time, formalized according to reference time ontology (e.g., OWL-time[17]);
– location information (e.g., "dbpedia:New_York_City"), grounding the event geographically.

This output perfectly fits the 4-feature ontologically necessary characterization of events introduced at the beginning of Section 2.

4 Ontological-based reasoning tasks for events

Provided the formal ontological characterization of event types presented in the analysis in Section 2, our goal is to demonstrate that this knowledge enables us to deal with different high-level reasoning tasks on events and their properties. In the following, we briefly exemplify some of these possible tasks and, following our approach, we mention what kind of information from the ontological analysis can be used; in the remainder of the paper, we will then turn our attention to the specific task of event coreference.

Pre- and Post-condition inference. One task regards the ability to reason on the states that enable the activation of an event (i.e. its pre-conditions), the effects on the knowledge that the event produces on the state (i.e. its post-conditions), and possibly the facts that are true during the event execution (aka "implied situations" in [22]). For example, supposing that we extract the event *"John has been hired by Google"*, then we would like to be able to infer that before this event *"John was not working in Google"* while later this fact is true.

The ontological properties that can be used for this task are clearly the information about status before, during and after events and possibly information on super- and sub-events with their associated knowledge content.

Completion of missing events. Another reasoning task can regard the inference of missing or implicit events and the completion of the events set on the base of the available event information. For example, suppose that we detect the two events *"Jobs is the CEO of Apple"*, and *"Cook is the CEO of Apple"* at different dates. We would like to infer that an event of the appointing of the new CEO of Apple has happened between the two events.

In this case, the useful ontological features can regard the status before, during and after the event, the possible asserted relations across event types (e.g. causality) and the local information about the participants.

Event information refinement. On the base of the ontological information, we might also want to refine the extracted information associated to an event in order to reach a more fine grained representation of the event local knowledge. This can be achieved possibly using contextual information known from the previously recognized events or background knowledge. For instance, suppose that we extract the event *"The president visited the capital of France and met the prime minister"* in the time period of July 2016 and American politics news, then we might infer that *"Barack Obama visited Paris and met Manuel Valls"*.

This task requires ontological information about roles and the ability to reason with contextual information (e.g. on the time and location of the event occurence).

Incompatible events. From the text extraction, we can obtain two events that have incompatible information (e.g. possibly coming from different sources): we want to be able to find that the

[12] http://www.framebase.org/
[13] http://www.adampease.org/OP/
[14] http://framenet.icsi.berkeley.edu/
[15] https://verbs.colorado.edu/~mpalmer/projects/verbnet.html
[16] http://dbpedia.org/
[17] https://www.w3.org/2006/time

happening is the same but the information is incompatible (due to errors or simply different versions of the facts). For example, we might have that a news reports *"Last night, John Smith killed two people with a knife"* while another news might report *"Last night, an unidentified suspect killed a woman with a knife"*: without knowing which version of the facts actually happened, we want to be able to assert that the two descriptions are incompatible.

To be able to obtain these inferences, we might want to compare the event descriptions that refer to the same event with respect to their roles, time and location in order to recognize the incompatible information and use information about incompatible states.

5 Event coreference task

In the following, we will first provide a definition of event coreference, then we show how the task of coreference discovery has been formalized and implemented in our experiments on the base of the ontological representation of events.

5.1 Coreference definition

In this paper, we define *event coreference* as an equivalence relation $\equiv_C$ among events: $e_1 \equiv_C e_2$ holds when events e_1 and e_2 describe the same happening or parts of the same happening regarding the four characterizing features described in Section 2. In other terms, the features of e_1 and e_2 identify different classifications (perspectives) of the same happening or the set of features of e_1 (e_2) identifies a part of the happening identified by e_2 (e_2, resp.ly).

Note that even when event coreference holds between two events e_1 and e_2, these events can still differ in terms of information content. For instance, if the event is an olympic race, e_1 may have information about the weather during the race which might not be available in event e_2. Indeed, for event coreference to be satisfied it suffices that they both have the same information regarding the event type (a race), time, participants and location.

The task that we consider with respect to event coreference is to determine the existence and degree of truth to which a coreference relation exists based on the related data on the events. Our underlying hypothesis is that an analysis at the event level, based on the proposed ontological properties and linguistic resources (e.g. FrameNet [2]), can emend and complete the information extracted from single event detection.

5.2 Coreference measure and aggregation

We can formally define a coreference measure and a relation to aggregate event descriptions as follows.

We call a *coreference measure* μ_t for event type t any function that assigns to every pair of event descriptions ed_1, ed_2 of type t a value in $[0, 1]$. Intuitively a coreference measure, provides an estimation of the plausibility for two event descriptions to refer to the same event.

A *coreference aggregation* α_t for event type t is a function that given two event descriptions, ed_1, ed_2 of type t returns a new joint event description ed_{1+2}. A coreference aggregation provides a way to merge two event descriptions into a single event description that aggregates all the information contained in the starting event descriptions.

5.3 Implementing coreference conditions rules

In our approach, rules for coreference measure and aggregation are formulated starting from the ontological definition of events.

In our first formulation for coreference rules we only considered the conditions for coreference measure: that is, in our experiments we are not interested in aggregating event descriptions but we only recognize and mark them as (possibly) coreferring.

As a first formulation for coreference measure, we simply provided rules for *certain* and *possible* coreference: the first conditions correspond to sufficient and necessary conditions (for each event type) for coreference existence, while the latter constitute only necessary conditions, providing an

evidence for possible existence of coreference. Such conditions have been extracted from the ontological analysis of the considered event type: in other words, these coreference criteria correspond to necessary (and sometime sufficient) event identity criteria that result from that analysis.

In the formulation of the rules, we considered the availability of FrameNet frame elements data and the event type properties defined in the ontological analysis. These rules have then been implemented as SPARQL rules and used in the evaluation of our approach, as described in next section.

The complete set of rules can be found in Appendix B. In Table 2, for example, we report the conditions we identified for the event type *Killing*. In the table, each row represents one condition: E_1 and E_2 are variables identifying two *Killing* event instances and we use the dot notation to denote the value of their properties (viz. frame elements). With $==$ we indicate equality of the values and with $\sim$ their compatibility (e.g. time value of E_1 can be included of the time value for E_2). The relation $hasCoref$ indicates that the compared (sub-)events are themselves recognized as coreferring.

Table 2. Conditions for event type *Killing*

certain coref.	$(E1.Victim == E2.Victim)$
	$(E1.SubEvent\ hasCoref\ E2.SubEvent)$
possible coref.	$(E1.Killer == E2.Killer)\ \&\ (E1.Time \sim E2.Time)$
	$(E1.Killer == E2.Killer)\ \&\ (E1.Place \sim E2.Place)$
	$(E1.Killer == E2.Killer)\ \&\ (E1.Tool == E2.Tool)$
	$(E1.Tool == E2.Tool)\ \&\ (E1.Time \sim E2.Time)$
	$(E1.Place \sim E2.Place)\ \&\ (E1.Time \sim E2.Time)$

We remark that in this version of the rules we only considered "intra-type" coreference rules: this choice was made in order to compare with the existing coreference approaches in the evaluation. However, it is easy to extend the set of rules with conditions establishing coreference relations across events of different types. For example, considering *Killing* and *Dying* events, we can assert that they (certainly) describe the same happening if the victim is the same individual: we can write this condition as

$$(E1.Victim == E2.Protagonist)\ \&\ Killing(E1)\ \&\ Dying(E2)$$

6 Evaluation on ECB benchmark

In this section, we evaluate the extraction of the coreference on the extended Event Coreference Bank corpus [3]. The dataset is a collection of 482 text documents from Google News, divided into 43 sets representing different topics. For each topic, the set of texts is annotated with intra- and cross-document event coreference, in accordance with the TimeML specification.

In literature, there are different ways to evaluate event coreference, each of them having different peculiarity: e.g., MUC [23], B^3 [1], CEAF [14], and BLANC [15]. BLANC, in particular, was developed to overcome the limitations of the other previously proposed metrics, as it takes into consideration both coreference and non-coreference links, does not ignore singletons (problematic for MUC) and does not boost the score in their presence (as typically occurs for B^3 and CEAF scores).

To compute these metrics, we relied on an already available package, CorScorer[18], an open source tool that measure coreference sets w.r.t. the most important metrics used for this kind of evaluation.

[18] `http://conll.github.io/reference-coreference-scorers/`

6.1 Experiments setup

Unfortunately, not every event in the ECB corpus is annotated in the document, therefore we choose to consider, in our classification, only the words that are actually annotated in the gold standard.

In addition, as we are currently dealing with the most-frequent 6 event types, we filter out the lemmas that do not appear as lexical units in the corresponding frames in FrameNet. In particular, our evaluation is restricted to 137 different lemmas.

We develop a lemma-based baseline to be compared with the results obtained by our system: all predicates in FrameNet that have the same lemma are coreferential.

We processed the ECB corpus with Pikes, obtaining an RDF knowledge graph of the knowledge about events, entities, locations, etc., conveyed by the news documents in the corpus. We then applied the rules expressing coreference conditions on this RDF knowledge graph. The rules were implemented as SPARQL based rules, enriching the input knowledge graph with additional `possibleCoref` or `certainCoref` triples based on the condition expressed in the body. The rules were applied to the RDF data using RDFpro [7].

6.2 Experiments results

Results of our experiments are shown in Table 3. Besides the baseline, we report the performances of applying possible and certain event coreference rules separately, and altogether. A total of 152 possible event coreference pairs and 275 certain event coreference pairs were generated. On BLANC,

Table 3. Experiments Results (numbers are percentages)

	MUC			B^3			CEAF (M)			BLANC		
	p	r	F1	p	r	F1	p	r	F1	p	r	F1
lemma baseline	81.78	89.14	85.30	39.87	73.46	51.69	45.48	45.48	45.48	62.45	76.08	66.14
only certain	100	7.39	13.77	100	2.76	5.38	86.34	6.96	12.88	96.47	88.38	91.6
only possible	100	6.08	11.46	100	3.35	6.5	82.14	5.84	10.91	95.07	90.54	92.18
possible + certain	100	13.48	23.75	100	6.12	11.54	84.37	12.8	22.23	98.23	89.6	93.29

our approach clearly outperforms the lemma baseline (approx. 27% difference on F1, considering both probable and certain coreference). On all metrics, our work consistently shows better precision scores than the baseline, while the contrary holds for recall. The good performances on precision should not surprise, as our rules get fired only when several conditions on the event description are met, and works also on events of the same type but expressed with different lemmas. On the contrary, the low performances on recall, especially compared with the baselines, are partly justified by the nature of the dataset considered in the evaluation: it consists of 43 different topics, and within each topic all documents mostly refer to the very same event, thus an approach based on matching lemma predicates it is likely to show good overall performances. However, in a general event extraction context, especially when dealing with event extraction from large news documents, such baseline approach will likely corefer very different event mentions just because they share the same predicate. In such context, precision-tuned approaches, as our work, tend to be more successful as they better cope with the noise and redundancy of information. As a final remark, our approach performs better when considering both certain and possible rules.

7 Related work

Determining when two event descriptions in a text corpus are about the same event is a challenging, popular research stream. Several approaches have been proposed, and we summarize the main ones in this section. Approaches typically exploit the event arguments (e.g., action, participants, location, time) when comparing and clustering events descriptions (cfr. [24]).

Among them, some approaches [10,6] adopt models consisting of labeled training data. [10] proposed an ontology-based deterministic clustering algorithm to group event descriptions of certain type. [6] presented a learning-based classification algorithm exploiting event arguments and features.

In [3] an unsupervised Bayesian clustering model is proposed for within- and cross-document event coreference. It is based on the hierarchical Dirichlet process, and grounded on the assumption that two event descriptions corefer if they have the same event properties and share the same event participants.

[12] presented an iterative approach, where entity coreference and event coreference are jointly performed, so that at each iteration, the intermediate results produced for one of the two tasks can be used in solving the other.

In [13] a supervised method for event coreference is proposed: the method exploits a rich feature set, and propagates information (alternatively) between events and their arguments.

Finally, [9] considered granularity (e.g., durations of event actions, granularity level of locations) in computing event coreference. The authors present a supervised pairwise binary classifier based on decision-trees which computes the compatibility of event attributes (e.g., event trigger, time, location, human and non-human participants).

Notwithstanding some attempts to use ontologies (e.g. [10]), none of the state-of-the-art approaches for within- and cross-document event coreference makes use of an ontological representation of events as presented in this paper.

On the other side, there is a long history of proposals for ontological definitions of events, for example inside DOLCE [4], conceptualization efforts as the *Event Model F* [21] or ontology design patterns[19]: in the context of Semantic Web data, some works considered reasoning with ontologically well founded definitions of events in a noisy and incomplete knowledge scenario like the case of (text extracted) news data. For instance, [16] introduces a formalization of composite events through a formal ontology called *Ontology of Complex Events (OntoCE)*. This work aims at integrating complex event information from different media by managing incompatibilities across their contents. Such definition of events is then used to reason for checking temporal consistency of complex events, find new temporal relations and perform causal reasoning. On the other hand, in events extraction, ontological models (like the ESO ontology [22]) are generally used to structure and expose events objects and reasoning is only applied at a high level, assuming already clean and reconciled representations.

Other works do not aim at a foundational characterization of event types, but concentrate on defining the properties of events in particular domains and (textual) resources. For example, [5] focus on characterizing text extracted events for RDF representation and reasoning over archeological procedures. Similarly, [20] outlines a project of an ontology-based framework to formally represent and extract historical domain events from archives in form of linked data: purpose of this extraction is to semi-automatically build narratives of complex events.

In the area of event representation, we may also cite works regarding Semantic Web enabled event stream reasoning: for example, [11] surveys the features of different approaches for analyzing streams of social media data to extract complex events. The authors, however, notice that the current approaches are limited by non-realistic assumptions of well-structured data, known streams and static ontologies. Current approaches lack support for temporal and spatial reasoning on complex events and usually do not support RDFS or OWL reasoning, restricting the possibilities to use background knowledge (thus including ontological definitions of events).

As noticed, the proposed approaches are usually interested in reasoning on high level (ontological) features and classification of events: in our work, our goal is to demonstrate that also "low level" processing tasks as event coreference can benefit from a well-founded characterization of events. Moreover, the contribution of our work does not reside in showing the benefits of using an event ontology for the execution of such tasks, but in the application of a prior ontological analysis of events types present in the resources.

8 Conclusion

In this paper we presented a method to derive coreference of text-extracted events on the base of a well-founded ontological analysis of events and their relations. We first provided a formal definition of event and presented a form of ontological analysis of event types. On the base of this, we defined a set of conditions for the discovery of event coreference, which we easily implemented as RDF-based reasoning rules. Finally, we evaluated our approach and hypotheses over the ECB corpus:

[19] See, for example, `http://ontologydesignpatterns.org/wiki/Submissions:EventCore`

the experiments show promising results encouraging us for further investigation in the direction of this initial work.

In this regard, we plan on extending the current work by considering a larger set of event types to be analyzed and the possiblity to apply our method to larger news datasets. Moreover, we also want to apply the method to different event reasoning tasks, like the ones identified in Section 4.

Appendix A Analysis of Event Types

A.1 Arresting

Ontological classification	Achievement
Active participant	Person (role: public officer)
Passive participant	Person (role: civil agent)
Tool participant	Law
Subevents	–
Sovraevents	Declaration of arrest
Status: before	Passive participant is alive
Status: during	Non applicable
Status: after	Passive participant is alive and subject to (legal) restrictions
Relationship among events	Arrest causes: being into legal custody (necessary); being prisoner (necessary); being handcuffed (possible)
Symmetrical events	To give up rights to the legal authority
Incompatible events/state	Events in which an active or passive participant of this event is released, dead or unconscious; a simultaneous arrest event in which the passive participant of this event is an active participant of the other (and vice versa)
Number of participants	Arrest must have: one or more active participants; one or more passive participants; one or more tool participants
Spatial region	Arrest happens in the location where the active and passive participants are located
Temporal region	Atomic
Repeatability	There can be several arrest events with the same passive participants; there can be more than one arrest event with the same active participants; there can be more than one arrest event in the same spatio-temporal region

A.2 Killing

Ontological classification	Accomplishment
Active participant	Physical object (including agents)
Passive participant	Living entity
Tool participant	Physical object (including chemical and biological entities)
Subevents	Hitting (possible); start of the killing, dying (necessary)
Sovraevents	Murdering, starving, being infected (all possible)
Status: before	Passive participant is alive
Status: during	Passive participant is alive
Status: after	Passive participant is dead
Relationship among events	Killing causes: damaging
Symmetrical events	Dying for causes external to the passive participant
Incompatible events/state	–
Number of participants	Killing can have one or more active participants; one or more passive participants; one or more tool participants.
Spatial region	Killing happens in the location where the passive participants are located
Temporal region	Interval
Repeatability	There cannot be two killing events with the same passive participant; there can be more than one killing event with the same active participant; there can be more than one killing event in the same spatio-temporal region.

A.3 Dying

Ontological classification	Achievement
Active participant	–
Passive participant	Living entity
Tool participant	Physical object (including chemical and biological entities)
Subevents	–
Sovraevents	Killing
Status: before	Passive participant is alive
Status: during	Non applicable
Status: after	Passive participant is dead
Relationship among events	Dying is caused by: murdering (possible); being infected (possible); begin poisoned (possible)
Symmetrical events	–
Incompatible events/state	Being born; acting
Number of participants	Dying can have one or more passive participants; one or more tool participants
Spatial region	Dying happens in the region where the passive participant is located
Temporal region	Atomic
Repeatability	There can be only one dying event with the same passive participant; there can be only one dying event in the same spatiotemporal region.

A.4 Charging (via legal notification)

Ontological classification	Achievement
Active participant	Person (role: public officer)
Passive participant	Person (role: civil agent)
Tool participant	Law, written text (depending on the law)
Subevents	–
Sovraevents	–
Status: before	Passive participant is alive and conscious, a judge rules the charge against the passive participant
Status: during	Non applicable
Status: after	The passive participant is charged
Relationship among events	Charging causes: acquisition of a legal status (necessary); begin under spacial legal restrictions (possible); being arrested (possible)
Symmetrical events	–
Incompatible events/state	Being dead; being unconscious
Number of participants	Charging has one or more active participants; one or more passive participants, one or more tool participants
Spatial region	Charging happens in the region where the passive participant is located
Temporal region	Atomic
Repeatability	There can be several charging events with the same active and passive participants; there can be several charging events in the same spatiotemporal region.

A.5 Shooting

Ontological classification	Accomplishment
Active participant	Agent, physical object (gun, bow)
Passive participant	Physical object (bullet, arrow, stone), amount of matter (sand, water)
Tool participant	Physical object (depending on active and passive participants: gun for person and bullet; bow for person and arrow, non applicable for person and spear)
Subevents	Exercising force, throwing, moving
Sovraevents	Murdering, attacking
Status: before	Active participant controls the passive participant or the tool participant, the tool participant is loaded
Status: during	Active participant controls the passive participant or the tool participant, the tool participant is unloaded
Status: after	Passive participant is moving, tool participant is unloaded (possible)
Relationship among events	Shooting causes: moving of the passive participant (necessary)
Symmetrical events	–
Incompatible events/state	Staying still, having no propelling power, being unloaded
Number of participants	Shooting can have one or more active participants; one or more passive participants; one or more tool participants
Spatial region	Shooting happens in the region where the active and passive participants are located
Temporal region	Interval
Repeatability	There can be several shooting events with the same passive participant; there can be several shooting events with the same active participant; there can be several shooting events in the same spatiotemporal region

A.6 Attacking

Ontological classification	Accomplishment
Active participant	Agent, biological entity (including system)
Passive participant	Agent, physical object, biological entity (including system)
Tool participant	Biological entity, physical object, amount of matter
Subevents	Start of the attack; end of the attack; hitting; destroying, moving
Sovraevents	Fighting, competing
Status: before	Active participant is alive or functioning, has control of the tool (if present)
Status: during	Active participant acts and controls the tool (if present)
Status: after	–
Relationship among events	Attacking causes: harming passive participant (possible); defending (possible)
Symmetrical events	–
Incompatible events/state	Being dead, being unconscious (for active participant)
Number of participants	Attacking can have one or more active participants; one or more passive participants; one or more tool participants
Spatial region	Attacking happens in the region where the tool and passive participants are located
Temporal region	Interval
Repeatability	there can be several attacking events with the same passive participant; there can be several attacking events with the same active participant; there can be several attacking events in the same spatiotemporal region

Appendix B Coreference conditions

B.1 Arresting

certain coref.	$(E1.Suspect == E2.Suspect)$ & $(E1.Time \sim E2.Time)$
possible coref.	$(E1.Suspect == E2.Suspect)$ & $(E1.Offense == E2.Offense)$
	$(E1.Suspect == E2.Suspect)$ & $(E1.Place \sim E2.Place)$
	$(E1.Place \sim E2.Place)$ & $(E1.Time \sim E2.Time)$ & $(E1.Offense == E2.Offense)$

B.2 Killing

certain coref.	$(E1.Victim == E2.Victim)$
	$(E1.SubEvent\ hasCoref\ E2.SubEvent)$
possible coref.	$(E1.Killer == E2.Killer)$ & $(E1.Time \sim E2.Time)$
	$(E1.Killer == E2.Killer)$ & $(E1.Place \sim E2.Place)$
	$(E1.Killer == E2.Killer)$ & $(E1.Tool == E2.Tool)$
	$(E1.Tool == E2.Tool)$ & $(E1.Time \sim E2.Time)$
	$(E1.Place \sim E2.Place)$ & $(E1.Time \sim E2.Time)$

B.3 Dying

certain coref.	$(E1.Protagonist == E2.Protagonist)$
possible coref.	$(E1.Place\ E2.Place)$ & $(E1.Time \sim E2.Time)$
	$(E1.Cause == E2.Cause)$ & $(E1.Time \sim E2.Time)$
	$(E1.Killer == E2.Killer)$ & $(E1.Tool == E2.Tool)$
	$(E1.Tool == E2.Tool)$ & $(E1.Time \sim E2.Time)$
	$(E1.Place \sim E2.Place)$ & $(E1.Time \sim E2.Time)$

B.4 Charging

certain coref.	$(E1.Accused == E2.Accused)$ & $(E1.Time == E2.Time)$
possible coref.	$(E1.Accused == E2.Accused)$ & $(E1.Charges == E2.Charges)$
	$(E1.Arraign_authority == E2.Arraign_authority)$ & $(E1.Place == E2.Place)$
	$(E1.Place \sim E2.Place)$ & $(E1.Time \sim E2.Time)$

B.5 Shooting

certain coref.	$(E1.Agent == E2.Agent)$ & $(E1.Goal == E2.Goal)$ & $(E1.Time == E2.Time)$
	$(E1.SubEvent\ hasCoref\ E2.SubEvent)$
possible coref.	$(E1.Place \sim E2.Place)$ & $(E1.Time \sim E2.Time)$
	$(E1.Agent == E2.Agent)$ & $(E1.Time == E2.Time)$
	$(E1.Agent == E2.Agent)$ & $(E1.Goal == E2.Goal)$
	$(E1.Projectile == E2.Projectile)$ & $(E1.Time == E2.Time)$

B.6 Attacking

certain coref.	$(E1.Assailant == E2.Assailant)$ & $(E1.Victim == E2.Victim)$ & $(E1.Time == E2.Time)$ $(E1.SubEvent\ hasCoref\ E2.SubEvent)$
possible coref.	$(E1.Place \sim E2.Place)$ & $(E1.Time \sim E2.Time)$ $(E1.Assailant == E2.Assailant)$ & $(E1.Time == E2.Time)$ $(E1.Victim == E2.Victim)$ & $(E1.Time == E2.Time)$ $(E1.Weapon == E2.Weapon)$ & $(E1.Time == E2.Time)$

References

1. Bagga, A., Baldwin, B.: Algorithms for scoring coreference chains. In: Int. Conf. on Language Resources and Evaluation Workshop on Linguistics Coreference. pp. 563–566 (1998)
2. Baker, C.F., Fillmore, C.J., Lowe, J.B.: The Berkeley FrameNet project. In: Int. Conf. on Computational linguistics-Volume 1. pp. 86–90. Association for Computational Linguistics (1998)
3. Bejan, C.A., Harabagiu, S.M.: Unsupervised event coreference resolution with rich linguistic features. In: ACL 2010. pp. 1412–1422 (2010)
4. Borgo, S., Masolo, C.: Foundational choices in DOLCE. In: Staab, S., Studer, R. (eds.) Handbook on Ontologies, pp. 361–381. International Handbooks on Information Systems, Springer (2009)
5. Byrne, K., Klein, E.: Automatic extraction of archaeological events from text. In: Computer Applications and Quantitative Methods in Archaeology (2010)
6. Chen, Z., Ji, H.: Event coreference resolution: Algorithm, feature impact and evaluation. In: Events in Emerging Text Types (eETTs 2009) Workshop (2009)
7. Corcoglioniti, F., Rospocher, M., Mostarda, M., Amadori, M.: Processing billions of rdf triples on a single machine using streaming and sorting. In: ACM SAC 2015. pp. 368–375. ACM (2015)
8. Corcoglioniti, F., Rospocher, M., Palmero Aprosio, A.: Frame-based ontology population with pikes. IEEE Transactions on Knowledge and Data Engineering 28(12), 3261–3275 (Dec 2016), http://dx.doi.org/10.1109/TKDE.2016.2602206
9. Cybulska, A., Vossen, P.: Bag of events approach to event conference resolution. supervised classification of event templates. International Journal of Computational Linguistics and Applications 6(2), 9–24 (2015)
10. Humphreys, K., Gaizauskas, R., Azzam, S.: Event coreference for information extraction. In: Workshop on Operational Factors in Practical, Robust Anaphora Resolution for Unrestricted Texts. pp. 75–81. ACL (1997)
11. Keskisärkkä, R., Blomqvist, E.: Semantic complex event processing for social media monitoring - a survey. In: SMILE 2013. CEUR workshop proceedings (2013)
12. Lee, H., Recasens, M., Chang, A., Sur-deanu, M., Jurafsky, D.: Joint entity and event coreference resolution across documents. In: EMNLP-CoNLL 2012. pp. 489–500. ACL (2012)
13. Liu, Z., Araki, J., Hovy, E., Mitamura, T.: Supervised within-document event coreference using information propagation. In: LREC 2014. pp. 4539–4544 (2014)
14. Luo, X.: On coreference resolution performance metrics. In: HLT 2005. pp. 25–32. ACL (2005)
15. Luo, X., Pradhan, S., Recasens, M., Hovy, E.: An extension of blanc to system mentions. In: Annual Meeting of ACL (Short Papers). pp. 24–29 (2014)
16. Mele, F., Sorgente, A.: The temporal representation and reasoning of complex events. In: CILC 2011. pp. 385–399 (2011)
17. Ovchinnikova, E., Vieu, L., Oltramari, A., Borgo, S., Alexandrov, T.: Data-driven and ontological analysis of framenet for natural language reasoning. In: LREC 2010. pp. 3157–64 (2010)
18. Pianesi, F., Varzi, A.: Events and event talk: An introduction. speaking of events (2000)
19. Rospocher, M., van Erp, M., Vossen, P., Fokkens, A., Aldabe, I., Rigau, G., Soroa, A., Ploeger, T., Bogaard, T.: Building event-centric knowledge graphs from news. Web Semantics: Science, Services and Agents on the World Wide Web 37–38, 132 – 151 (2016)
20. Rovera, M.: A knowledge-based framework for events representation and reuse from historical archives. In: ESWC 2016. pp. 845–852 (2016)
21. Scherp, A., Franz, T., Saathoff, C., Staab, S.: F–a model of events based on the foundational ontology DOLCE+DnS ultralight. In: K-CAP 2009. pp. 137–144. ACM (2009)
22. Segers, R., Vossen, P., Rospocher, M., Serafini, L., Laparra, E., Rigau, G.: ESO: a Frame based Ontology for Events and Implied Situations. In: MAPLEX 2015 (2015)
23. Vilain, M., Burger, J., Aberdeen, J., Connolly, D., Hirschman, L.: A model-theoretic coreference scoring scheme. In: MUC6 1995. pp. 45–52. ACL (1995)
24. Vossen, P., Rodrigo Agerri, I.A., Cybulska, A., van Erp, M., Fokkens, A., Laparra, E., Minard, A.L., Aprosio, A.P., Rigau, G., Rospocher, M., Segers, R.: Newsreader: Using knowledge resources in a cross-lingual reading machine to generate more knowledge from massive streams of news. Knowledge-Based Systems 110, 60 – 85 (2016), http://www.sciencedirect.com/science/article/pii/S0950705116302271